

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
20 February 2003 (20.02.2003)

PCT

(10) International Publication Number
WO 03/014911 A2

(51) International Patent Classification⁷:

G06F 3/06

(21) International Application Number:

PCT/US02/25291

(22) International Filing Date:

9 August 2002 (09.08.2002)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

09/928,128 10 August 2001 (10.08.2001) US

(71) Applicant: SUN MICROSYSTEMS, INC. [US/US];
4150 Network Circle, Santa Clara, CA 95054 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(72) Inventors: CARLSON, Mark, A.; 1021 Vivian Circle, Boulder, CO 80303 (US). DA SILVA, Rowan, E.; 11 Old Boston Road, #502, Tewksbury, MA 01876 (US).

Published:

— without international search report and to be republished upon receipt of that report

(74) Agents: KONRAD, William, K. et al.; Konrad Raynes & Victor LLP, Suite 210, 315 South Beverly Drive, Beverly Hills, CA 90212 (US).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD, SYSTEM, AND PROGRAM FOR GENERATING AND USING CONFIGURATION POLICIES

(57) Abstract: Provided is a method, system, and program for managing multiple resources in a system. A user request is received to generate a configuration policy. User selection is received of a set of the multiple resources. A determination is made of at least one element for each selected resource in the set, wherein each element is capable of managing one of the resources in the system. User selection is received of one element for each selected resource in the set. The configuration policy is defined to include the user selected elements, wherein invoking the configuration policy further invokes each element defined in the configuration policy to configure the resources associated with the invoked elements.

WO 03/014911 A2

METHOD, SYSTEM, AND PROGRAM FOR
GENERATING AND USING CONFIGURATION POLICIES
RELATED APPLICATIONS

[0001] This application is related to the co-pending and commonly assigned patent
5 application "Method, System, And Program For Managing Multiple Resources in a
System", by Mark A. Carlson, and having attorney docket no. P5445, which
application is filed on the same date herewith and incorporated herein by reference in
its entirety.

10

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0002] The present invention relates to a method, system, and program for managing
multiple resources in a system.

15

2. Description of the Related Art

[0003] A storage area network (SAN) comprises a network linking one or more
servers to one or more storage systems. Each storage system could comprise a
Redundant Array of Independent Disks (RAID) array, tape backup, tape library, CD-
ROM library, or JBOD (Just a Bunch of Disks) components. Storage area networks
20 (SAN) typically use the Fibre Channel protocol, which uses optical fibers to connect
devices and provide high bandwidth communication between the devices. In Fibre
Channel terms the one or more switches interconnecting the devices is called a
"fabric". However, SANs may also be implemented in alternative protocols, such as
InfiniBand**, IPStorage over Gigabit Ethernet, etc.

25

[0004] In the current art, to add or modify the allocation of storage or other
resources in a SAN, an administrator must separately utilize different software
programs to configure the SAN resources to reflect the modification to the storage
allocation. For instance to allow a host to alter the allocation of storage space in the
SAN, the administrator would have to perform one or more of the following:

30

- use a storage device configuration tool to resize a logical volume, such as a
logical unit number (LUN), or change the logical volume configuration at the

storage device, e.g., the RAID or JBOD, to provide more or less storage space to the host.

- a switch configuration tool to alter the assignment of paths in the switch to the host, i.e., rezoning, to provide access to the newly reconfigured logical volume (LUN).

5

- perform LUN masking, which involves altering the assignment of HBA interface ports to the reconfigured LUNs.

10

- use a host volume manager configuration tool to alter the allocation of physical storage to logical volumes used by the host. For instance if the administrator adds storage, then the logical volume must be updated to reflect the added storage.

15

- use a backup program manager to reflect the change in storage allocation so that the backup program will backup more or less data for the host.

20

- use a snapshot copy configuration manager to update the host logical volumes that are subject to a snapshot copy, where a backup copy is made by copying the pointers in the logical volume.

[0005] Not only does the administrator have to invoke one or more of the above tools to implement the requested storage allocation change throughout the SAN, but the administrator may also have to perform these configuration operations repeatedly 25 if the configuration of multiple distributed devices is involved. For instance, to add several gigabytes of storage to a host logical volume, the administrator may allocate storage space on different storage subsystems in the SAN, such as different RAID boxes. In such case, the administrator would have to separately invoke the configuration tool for each separate device involved in the new allocation. Further, 30 when allocating more storage space to a host logical volume, the administrator may have to allocate additional storage paths through separate switches that lead to the one

or more storage subsystems including the new allocated space. The complexity of the configuration operations the administrator must perform further increases as the number of managed components in a SAN increase. Moreover, the larger the SAN, the increased likelihood of hosts requesting storage space reallocations to reflect new storage allocation needs.

[0006] Additionally, many systems administrators are generalists and may not have the level of expertise to use a myriad of configuration tools to appropriately configure numerous different vendor resources. Still further, even if an administrator develops the skill and knowledge to optimally configure networks of components from different 10 vendors, there is a concern for knowledge retention in the event the skilled administrator separates from the organization. Yet further, if administrators are not utilizing their configuration knowledge and skills, then their skill level at performing the configurations may decline.

[0007] All these factors, including the increasing complexity of storage networks, 15 decreases the likelihood that the administrator may provide an optimal configuration.

[0008] The above described difficulties in configuring resources in a Fibre Channel SAN

environment are also experienced in other storage environments including multiple storage devices, hosts, and switches, such as InfiniBand**, IPStorage over Gigabit Ethernet, etc.

[0009] For all the above reasons, there is a need in the art for an improved technique for managing and configuring the allocation of resources in a large network, such as a SAN.

25 SUMMARY OF THE PREFERRED EMBODIMENTS

[0010] Provided is a method, system, and program for managing multiple resources in a system. A user request is received to generate a configuration policy. User selection is received of a set of the multiple resources. A determination is made of at least one element for each selected resource in the set, wherein each element is capable of managing one of the resources in the system. User selection is received of one element for each selected resource in the set. The configuration policy is defined

to include the user selected elements, wherein invoking the configuration policy further invokes each element defined in the configuration policy to configure the resources associated with the invoked elements.

[0011] In further implementations, the multiple resources include a storage device, a 5 switch, a host adaptor, and a volume manager. The element managing the storage device allocates the storage space to the host; the element managing the switch is capable of allocating at least one path in the switch to the storage device to allow the host to access the allocated storage space; the element managing the host adaptors allocates at least one host adaptor in the host to communicate with the switch to 10 access the allocated storage space; and the element managing the volume manager assigns the allocated storage space in the device to the requested logical volume used by the host.

[0012] Still further, each of multiple elements provided for one resource define a different configuration of the resource.

15 [0013] Further provided is a method, system, and program for configuring multiple resources in the system. User selection is received of one of multiple configuration policies, wherein each configuration policy defines resources to configure and one element for each resource to configure, and wherein each element specifies configuration parameters to use to configure the resource. User selection is further 20 received of an instance of one resource to configure, wherein the user selected resource instance is capable of being configured by the configuration policy. A determination is made of additional resource instances that are configured by the selected configuration policy. The elements defined for the selected configuration policy are then called to configure the user selected resource instance and the 25 determined additional resource instances according to the element configuration parameters.

[0014] In further implementations, one of the resources to configure comprises a storage device. In such case, a determination is made of available storage space at each storage device instance that is available to the user selected resource instance.

30 User selection is received of an amount of storage space to allocate and a determination is made of one storage device instance including the user selected

amount of storage space. Calling the elements to configure each user selected resource further comprises calling a storage element to configure the determined storage device instance to allocate the user selected amount of storage space to the configuration.

5

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] Referring now to the drawings in which like reference numbers represent corresponding parts throughout:

FIG. 1 illustrates a network computing environment for one implementation of
10 the invention;

FIG. 2 illustrates a component architecture in accordance with certain
implementations of the invention;

FIG. 3 illustrates a component architecture for a storage network in accordance
with certain implementations of the invention;

15 FIG. 4 illustrates logic to invoke a configuration operation in accordance with
certain implementations of the invention; and

FIG. 5 illustrates logic to configure network components in accordance with
certain implementations of the invention.

FIG. 6 illustrates further components within the administrator to define and
20 execute configuration policies in accordance with certain implementations of the
invention;

FIGs. 7-8 illustrate GUI panels through which a user invokes a configuration
policy to configure and allocate resources to provide storage space in accordance with
certain implementations of the invention; and

25 FIGs. 9-10 illustrate logic implemented in the configuration policy tool to
enable a user to invoke and use a defined configuration policy to allocate and
configure system resources in accordance with certain implementations of the
invention;

FIGs. 11-12 illustrate graphical user interface (GUI) panels through which a
30 user may define a configuration policy in accordance with certain implementations of
the invention; and

FIG. 13-14 illustrates logic implemented in a configuration policy tool to enable a user to create and define a configuration policy;

DETAILED DESCRIPTION

5 [0016] In the following description, reference is made to the accompanying drawings which form a part hereof and which illustrate several embodiments of the present invention. It is understood that other embodiments may be utilized and structural and operational changes may be made without departing from the scope of the present invention.

10 [0017] FIG. 1 illustrates an implementation of a Fibre Channel based storage area network (SAN) which may be configured using the implementations described herein. Host computers 4 and 6 may comprise any computer system that is capable of submitting an Input/Output (I/O) request, such as a workstation, desktop computer, server, mainframe, laptop computer, handheld computer, telephony device, etc. The 15 host computers 4 and 6 would submit I/O requests to storage devices 8 and 10. The storage devices 8 and 10 may comprise any storage device known in the art, such as a JBOD (just a bunch of disks), a RAID array, tape library, storage subsystem, etc. Switches 12a, b interconnect the attached devices 4, 6, 8, and 10. The fabric 14 comprises the switches 12a, b that enable the interconnection of the devices. In the 20 described implementations, the links 16a, b, c, d and 18a, b, c, d connecting the devices comprise Fibre Channel fabrics, Internet Protocol (IP) switches, Infiniband fabrics, or other hardware that implements protocols such as Fibre Channel Arbitrated Loop (FCAL), IP, Infiniband, etc. In alternative implementations, the different components of the system may comprise any network communication technology known in the art. Each device 4, 6, 8, and 10 includes multiple Fibre Channel 25 interfaces 20a, 20b, 22a, 22b, 24a, 24b, 26a, and 26b, where each interface, also referred to as a device or host bus adaptor (HBA), can have one or more ports. Moreover, actual SAN implementation may include additional storage devices, hosts, host bus adaptors, switches, etc., than those illustrated in FIG. 1.

30 [0018] A path, as that term is used herein, refers to all the components providing a connection from a host to a storage device. For instance, a path may comprise host

adaptor 20a, fiber 16a, switch 12a, fiber 18a, and device interface 24a, and the storage devices or disks being accessed.

[0019] Certain described implementations provide a configuration technique that allows administrators to select a specific service configuration policy providing the path availability, RAID level, etc., to use to allocate, e.g., modify, remove or add, storage resources used by a host 4, 6 in the SAN 2. After the service configuration policy is specified, the component architecture implementation described herein automatically configures all the SAN components to implement the requested allocation at the specified configuration quality without any further administrator involvement, thereby streamlining the SAN storage resource configuration and allocation process. The requested allocation of the configuration is referred to as a service configuration policy that implements a particular configuration requested by a by calling the elements to handle the resource configuration. The policy provides a definition of configurations and how these elements in SAN are to be configured. In certain described implementations, the configuration architecture utilizes the Sun Microsystems, Inc. ("SUN") Jiro distributed computing architecture.**

[0020] Jiro provides a set of program methods and interfaces to allow network users to locate, access, and share network resources, referred to as services. The services may include hardware devices, software devices, application programs, storage resources, communication channels, etc. Services are registered with a central lookup service server, which provides a repository of service proxies. A network participant may review the available services at the lookup service and access service proxy objects that enable the user to access the service through the service provider. A "proxy object" is an object that represents another object in another memory or program memory address space, such as a resource at a remote server; to enable access to that resource or object at the remote location. Network users may "lease" a service, and access the proxy object implementing the service for a period of time.

[0021] A service provider discovers lookup services and then registers service proxy objects and service attributes with the discovered lookup service. In Jiro, the service proxy object is written in the Java** programming language, and includes methods and interfaces to allow users to invoke and execute the service object located through

the lookup service. A client accesses a service proxy object by querying the lookup service. The service proxy object provides Java interfaces to enable the client to communicate with the service provider and access the service available through the network. In this way, the client uses the proxy object to communicate with the service provider to access the service.

[0022] FIG. 2 illustrates a configuration architecture 100 using Jiro components to configure resources available over a network 102, such as hosts, switches, storage devices, etc. The network 102 may comprise the fiber links provided through the fabric 14. The network 102 allows for communication among an administrator user 10 interface (UI) 104, one or more elements 106 (only one is shown, although multiple elements 106 may be present), one or more configuration policy services (only one is shown) 108, and a lookup service 110.

[0023] The network 102 may comprise the Internet, an Intranet, a LAN, etc., or any other network system known in the art, including wireless and non-wireless networks. 15 The administrator UI 104 comprises a system that submits requests for access to network resources. For instance, the administrator UI 104 may request a new allocation of storage resources to hosts 4, 6 (FIG. 1) in the SAN 2. The administrator UI 104 may be implemented as a program within the host 4, 6 involved in the new storage allocation or a within system remote to the host. The administrator UI 104 20 provides access to the configuration resources described herein to alter the configuration of storage resources to hosts. The elements 106 provide a management interface to provide configuration and control over a resource 112. In SAN implementations, the resource 112 may comprise any resource in the system that is configured during the process of allocating resources to a host. For instance, the 25 configurable resources 112 may include host bus adaptors 20a, b, 22a, b, a host volume manager which provides an assignment of logical volumes in the host 4, 6 to physical storage space in storage devices 8,10, a backup program in the host 4, 6, a snapshot program in the host 4, 6 providing snapshot services (i.e., copying of pointers to logical volumes), switches 12a, b, storage devices 8, 10, etc. Multiple 30 elements may be defined to provide different configuration qualities for a single resource. Each of the above components in the SAN would comprise a separate

resource 112 in the system, where one or more elements 106 are provided for management and configuration of the resource. The service configuration policy 108 implements a particular configuration requested by the host 104 by calling the elements 106 to configure the resources 112.

5 [0024] In the architecture 100, the element 106, service configuration policy 108, and resource APIs 126 function as Jiro service providers that make services available to any network participant, including to each other and to the administrator UI 104. The lookup service 110 provides a Jiro lookup service in a manner known in the art. The lookup service 110 maintains registered service objects 114, including a lookup 10 service proxy object 116, that enables network users, such as the administrator UI 104, elements 106, service configuration policies 108, and resource APIs 126 to access the lookup service 110 and the proxy objects 116, 118a...n, 119a...m, and 120 therein. For instance, each element 106 registers an element proxy object 118a..n, each resource API 126 registers an API proxy object 119a...m, and each service 15 configuration policy 108 registers a service configuration policy proxy object 120 to provide access to the underlying resources. The service configuration policy 108 includes code to call elements 106 to perform the user requested configuration operations to reallocate storage resources to a specified host and logical volume.

[0025] With respect to the elements 106, the resources 112 comprise the underlying 20 service resource being managed by the element 106, e.g., the storage devices 8, 10, host bus adaptors 16a, b, c, d, switches 12a, b, host volume manager, backup program, snapshot program, etc. The resource application program interfaces (APIs) 126 provide access to the configuration functions of the resource to perform the resource specific configuration operations. Thus, there is one resource API set 126 for each 25 managed resource 112. The APIs 126 are accessible through the API proxy objects 119a...m. Because there may be multiple elements to provide different configurations of a resource 112, the number of registered element proxy objects n may exceed the number of registered API proxy objects m , because the multiple elements 106 that provide different configurations of the same resource 112 would use the same set of 30 APIs 126.

[0026] The element 106 includes configuration policy parameters 124 that provide the settings and parameters to use when calling the APIs 126 to control the configuration of the resource 112. If there are multiple elements 106 for a single resource 112, then each of those elements 106 may provide a different set of configuration policy parameters 124 to configure the resource 112. For instance, if the resource 112 is a RAID storage device, then the configuration policy parameters 124 for one element may provide a RAID level abstract configuration, or some other defined RAID configuration, such as Online Analytical Processing (OLAP) RAID definitions and configurations which may define a RAID level, number of disks, etc.

5 10 Another element may provide a different RAID level. Additionally, if the resource 112 is a switch, then the configuration policy parameters 124 for one element 106 may configure redundant paths through the switch to the storage space to avoid a single point of failure, whereas another element for the switch may configure only a single path. Thus, the elements 106 utilize the configuration policy parameters 124 and the

15 15 resource API 126 to control the configuration of the resource 112, e.g., storage device 8, 10, switches 12a, b, volume manager, backup program, host bus adaptors (HBAs) 20a, b, 22a, b, etc.

[0027] Each service configuration policy 108 would call one of the elements 106 for each resource 112 to perform the administrator/user requested reconfiguration. There

20 20 may be multiple service configuration policies for different predefined configuration qualities. For instance, there may be a higher quality service configuration policy, such as "gold", for critical data that would call one element 106 for each resource 112 to reconfigure, where the called element 106 configures the resource 112 to provide for extra protection, such as a high RAID level, redundant paths through the switch to

25 25 the storage space to avoid a single point of failure, redundant use of host bus adaptors to further reduce a single point of failure at the host, etc. A "bronze" or lower quality service configuration policy may not require such redundancy and protection to provide storage space for less critical data. The "bronze" quality service configuration policy 108 would call the elements 106 that implement such a lower

30 30 quality configuration policy with respect to the resources 112. Each called element 106 in turn calls the APIs 126 for the resource to reconfigure. Note that different

service configuration policies 108 may call the same or different elements 106 to configure a particular resource.

[0028] Associated with each proxy object 118a..n, 119a...m, and 120 are service attributes 128a...n, 129a...n, and 130 that provide descriptive attributes of the proxy objects 118a..n, 119a...n, and 120. For instance, the administrator UI 104 may use the lookup service proxy object 116 to query the service attributes 130 of the service configuration policy 108 to determine the quality of service provided by the configuration policy, e.g., the RAID level, number of redundant paths, etc. The service attributes 128a...n for the elements 106 may describe the type of configuration 10 performed by the specific element.

[0029] FIG. 2 further illustrates a topology database 140 which provides information on the topology of all the resources in the system, i.e., the connections between the host bus adaptors, switches and storage devices. The topology database 140 may be created during system initialization and updated whenever changes are made to the 15 system configuration in a manner known in the art. For instance, the Fibre Channel and SCSI protocols provide protocols for discovering all of the components or nodes in the system and their connections to other components. Alternatively, out-of-band discovery techniques could utilize Simple Network Management Protocol (SNMP) commands to discover all the devices and their topology. The result of the discovery 20 process is the topology database 140 that includes entries identifying the resources in each path in the system. Any particular resource may be available in multiple paths. For instance, a switch may be in multiple entries as the switch may provide multiple paths between different host bus adaptors and storage devices. The topology database 140 can be used to determine whether particular devices, e.g., host bus adaptors, 25 switches and storage devices, can be used, i.e., are actually interconnected. The lookup service 114 maintains a topology proxy object 142 that provides methods for accessing the topology database 140 to determine how components in the system are connected.

[0030] When the service configuration policy proxy object 120 is created, the 30 topology database 140 may be queried to determine those resources that can be used by the service configuration policy 108, i.e., those resources that when combined can

satisfy the configuration policy parameters 124 of the elements 106 defined for the service configuration policy 108. The service configuration policy proxy object service attributes 130 may be updated to indicate the query results of those resources in the system that can be used with the configuration. The service attributes 130 may

5 further provide topology information indicating how the resources, e.g., host bus adaptors, switches, and storage devices, are connected or form paths. In this way, the configuration policy proxy object service attributes 130 defines all paths of resources that satisfy the configuration policy parameters 124 of the elements 106 included in the service configuration policy.

10 [0031] In the architecture of FIG. 2, the service providers 108 (configuration policy service), 106 (element), and resource APIs 126 function as clients when downloading the lookup service proxy object 116 from the lookup service 110 and when invoking lookup service proxy object 116 methods and interfaces to register their respective service proxy objects 118a...n, 119a...m, and 120 with the lookup service 110. The

15 client 104 and service providers 106 and 108 would execute methods and interfaces in the service proxy objects 118a...n, 119a...m, and 120 to communicate with the service provider 106, 108, and 126 to access the associated service. The registered service objects 118a...n, 119a...m, and 120 comprise the services available through the lookup service 110. The administrator UI 104 uses the lookup service proxy object 116 to

20 access the proxy objects from the lookup service 110. Further details on how clients may discover and download the lookup service and service objects and register service objects are described in the Sun Microsystem, Inc. publications: "Jini Architecture Specification" (Copyright 2000, Sun Microsystems, Inc.) and "Jini Technology Core Platform Specification" (Copyright 2000, Sun Microsystems, Inc.), both of which

25 publications are incorporated herein by reference in their entirety.

[0032] The resources 112, elements 106, service configuration policy 108, and resource APIs 126 may be implemented in any computational device known in the art and each include a Java Virtual Machine (JVM) and a Jiro package (not shown). The Jiro package includes all the Java methods and interfaces needed to implement the

30 Jiro network environment in a manner known in the art. The JVM translates methods and interfaces of the Jiro package as well as the methods and interfaces of

downloaded service objects, into bytecodes capable of executing on the configuration policy service 108, administrator UI 104 element 106, and resource APIs 126. Each component 104, 106, 108, and 110 further includes a network protocol stack (not shown) to enable communication over the network. The network protocol stack 5 provides a network address for the components 104, 106, 108, 110, and 126, such as a Transmission Control Protocol/Internet Protocol (TCP/IP) address, support for unicast and multicast broadcasting, and a mechanism to facilitate the downloading of Java files. The network protocol stack may also include the communication infrastructure to allow objects, including proxy objects, on the systems to communicate, such as the 10 Common Object Request Broker Architecture (CORBA), Remote Method Invocation (RMI), TCP/IP, etc.

[0033] As discussed, the configuration architecture may include multiple elements for the different configurable resources in the storage system. Following are the resources that may be configured through the proxy objects in the SAN:

15 Storage Devices: There may be a separate element service for each configurable storage device 8, 10. In such case, the resource 112 would comprise the configurable storage space of the storage devices 8, 10 and the element 106 would comprise the configuration software for managing and configuring the storage devices 8, 10 according to the configuration policy 20 parameters 124. The element 106 would call the resource APIs 126 to access the functions of the storage configuration software.

25 Switch: There may be a separate element service for each configurable switch 12a, b. In such case, the resource 112 would comprise the paths in the switch and the element 106 would comprise the switch software for managing and configuring paths within the switch 12a, b according to the configuration policy parameters 124. The element 106 would call the resource APIs 126 to access the functions of the switch configuration software.

30 Host Bus Adaptors: There may be a separate element service to manage the allocation of the host bus adaptors 20a, b, 22a, b on each host 4, 6. In such case, the resource 112 would comprise all the host bus adaptors (HBAs) on a given host and the elements would comprise the configuration software for

assigning the host bus adaptors (HBAs) to a path according to the configuration policy parameters 124. The element 106 would call the resource APIs 126 to access the functions of the host adaptor configuration software on each host 4, 6.

5 Volume Manager on the Host: There may be a separate element service for the volume manager on each host 4, 6. In such case, the resource 112 would comprise the mapping of logical to physical storage and the element 106 would comprise the software for configuring the mapping of the logical volumes viewed by the host 4, 6 to physical storage space in the storage devices 8, 10 according to the configuration policy parameters 124. The element 106 would call the resource APIs 126 to access the functions of the volume manager configuration software.

10 Backup Program on the Host: There may be a separate element service 106 for the backup program configuration at each host 4, 6. In such case, the resource 112 would comprise the configurable backup program for the host 4, 6 and the element 106 would comprise software for managing and configuring backup operations for the host 4, 6 according to the configuration policy parameters 124. The element 106 would call the resource APIs 126 to access the functions of the backup management software.

15 Snapshot on the Host: There may be a separate element service 106 for the snapshot configuration for each host 4, 6. In such case, the resource 112 would comprise the snapshot operation on the host and the element 106 would comprise the software to select logical volumes to copy as part of a snapshot operation according to the configuration policy parameters 124. The element 106 would call the resource APIs 126 to access the functions of the snapshot configuration software.

[0034] Element services may also be provided for other network based storage devices and host based storage software other than those described herein.

[0035] FIG. 3 illustrates an additional arrangement of the element, service configuration policies, and APIs for the SAN components that may be available over a network 200, including a gold 202 and bronze 204 quality service configuration polices, each providing a different quality of configuration for the system components.

- 5 The service configuration policies 202 and 204 call one device configuration element for each resource that needs to be configured. The component architecture includes one or more storage device element configurations 214a, b, c, switch element configurations 216a, b, c, host bus adaptor (HBA) element configurations 218a, b, c, and volume manager element configurations 220a, b, c. The configuration elements
- 10 214a, b, c, 216a, b, c, 218a, b, c, and 220a, b, c call the resource APIs 222, 224, 226, and 228, respectively, that enable access and control to the commands and functions used to configure the storage device 230, switch 232, host bus adaptors (HBA) 234, and volume manager 236, respectively. In certain implementations, the resource API proxy objects include service attributes that describe the availability of resources for
- 15 the device which the particular API resources manage, i.e., available storage space, available paths, available host bus adaptor, etc. In the described implementations, there is a separate resource API object for each instance of the device, such that if there are two storage devices in the system, then there would be two storage configuration APIs, each providing the APIs to one of the storage devices. Further,
- 20 the proxy object for each resource API would include service attributes describing the availability at the resource to which the resource API provides access.

[0036] Each of the service configuration policies 202 and 204, configuration elements 214a, b, c, 216a, b, c, 218a,b , c, and 220a, b, c, and resource APIs 222, 224, 226, and 228 would register their respective proxy objects with the lookup service

- 25 250. For instance, the service configuration policy proxy objects 238 include the proxy objects for the gold 202 and bronze 200 quality service configuration polices; the element configuration proxy objects 240 include the proxy objects for each element 214a, b, c, 216a, b, c, 218a, b, c, 220a, b, c configuring a resource 230, 232, 234, and 236; and the API proxy objects 242 include the proxy objects for each set of
- 30 device APIs 222, 224, 226, and 228. As discussed each service configuration policy 200, 202 would call one element for each of the resources 230, 232, 234, and 236 that

need to be configured to implement the user requested configuration quality. Each device configuration element 214a, b, c, 216a, b, c, 218a, b, c, and 220a, b, c maintains configuration policy parameters (not shown) that provides a particular quality of configuration of the managed resource. Moreover, additional device 5 element configurations would be provided for each additional devices in the system. For instance, if there were two storage devices in the SAN system, such as a RAID box and a tape drive, there would be separate element configurations to manage each different storage device and separate proxy objects and accompanying APIs to allow access to each of the element configurations for the storage devices. Further, there 10 would be one or more host bus adaptor (HBA) element configurations for each host system to allow configuration and management of all the host bus adaptors (HBAs) in a particular host 4, 6 (FIG. 1). Each proxy object would include service attributes providing information on the resource being managed, such as the amount of available disk space, available paths in the switch, available host bus adaptors at the host, 15 configuration quality and configuration parameters, etc.

[0037] An administrator user interface (UI) 252 operates as a Jiro client and provides a user interface to enable access to the lookup service proxy object 254 from the lookup service 250 and enable access to the lookup service proxy object 254 to access the service configuration policies 202 and 204. The administrator 252 is a 20 process running on any system, including the device components shown in FIG. 3, that provides a user interface to access, run, and modify configuration policies. The service configuration policies 202, 204 call the configuration elements 214a, b, c, 216a, b, c, 218a, b, c, and 220a, b, c to configure each resource 230, 232, 234, 236 to implement the allocation of the additional requested storage space to the host. The 25 service configuration policies 202, 204 would provide a graphical user interface (GUI) to enable the administrator to enter resources to configure. Before a user at the administrator UI 252 could utilize the above described component architecture of FIG. 3 to configure components of a SAN system, e.g., the SAN 2 in FIG. 1, the service configuration policies 202, 204, element configurations 214a, b, c, 216a, b, c, 218a, b, c, and 220a, b, c would have to discover and join the lookup service 250 to 30 register their proxy objects. Further, each of the service configuration policies 202

and 204 must download the element configuration proxy objects 240 for the elements 214a, b, c, 216a, b, c, 218a, b, c, and 220a, b, c. The elements 214a, b, c, 216a, b, c, 218a, b, c, and 220a, b, c, in turn, must download one of the API proxy objects 242 for resource APIs 222, 224, 226, and 228, respectively, to perform the desired

5 configuration according to the configuration policy parameters maintained in the element and the host storage allocation request.

[0038] FIG. 3 further shows a topology database 256 and topology proxy object 258 that maintains the topology information on the database. Each record may specify the resources in a path.

10 [0039] FIG. 4 illustrates logic implemented within the administrator UI 252 to begin the configuration process utilizing the configuration architecture described with respect to FIGs. 2 and 3. Control begins at block 300 with the administrator UI 252 (“admin”) discovering the lookup service 250 and downloading the lookup service proxy object 254. The administrator UI 252 then uses (at block 302) the interfaces of 15 the lookup service proxy object 254 to access information on the service attributes providing information on each service configuration policy 202 and 204, such as the quality of availability and path redundancy. A user may then select one of the service configuration policies 202 and 204 appropriate to the availability and redundancy needs of the application that will use the new allocation of storage. For instance, a 20 critical database application would require high availability and redundancy, whereas an application involving non-critical data requires less availability and redundancy. The administrator UI 252 then receives user selection (at block 304) of one of the service configuration policies 202, 204 and a host and logical volume and other device components, such as switch 232 and storage device 230 to configure for the new 25 storage allocation. The administrator UI 252 may execute within the host to which the new storage space will be allocated or be remote to the host.

[0040] The administrator UI 252 then uses (at block 306) interfaces from the lookup service proxy object 254 to access and download the selected service configuration policy proxy object. The administrator UI 252 uses (at block 308) 30 interfaces from the downloaded service configuration policy proxy object to

communicate with the selected service configuration policy 202 or 204 to implement the requested storage allocation for the specified logical volume and host.

[0041] FIG. 5 illustrates logic implemented in the service configuration policy 202, 204 and element configurations 214a, b, c, 216a, b, c, 218a, b, c, 220a, b, c to perform the requested configuration operation. Control begins at block 350 when the service configuration policy 202, 204 receives a request from the administrator UI 252 for a new allocation of storage space for a logical volume and host through the configuration policy service proxy object 238, 240. In response, the selected service configuration policy 202, 204 calls (at block 352) one associated element configuration proxy object for each resource 222, 224, 226, 228 that needs to be configured to implement the allocation. In the logic described at blocks 354 to 370, the service configuration policy 202, 204 configures the following resources, the storage device 230, switch 232, host bus adaptors 234, and volume manager 236 to carry out the requested allocation. Additionally, the service configuration policy 202, 204 may call elements to configure more or less resources. For instance, for certain configurations, it may not be necessary to assign an additional path to the storage device for the added space. In such case, the service configuration policy 202, 204 would only need to call the storage device element configuration 214a, b, c and volume manager element configuration 220a, b, c to implement the requested allocation.

[0042] At block 354, the called storage device element configuration 214a, b, c uses interfaces in the lookup service proxy object 254 to query the service attributes of the storage configuration APIs 222 for storage devices 230 in the system to determine one or more storage configuration API proxy objects capable of configuring storage device(s) 230 having enough available space to fulfill requested storage allocation with a storage type level that satisfies the element configuration policy parameters. For instance, the gold service configuration policy 202 will call device element configurations that provide for redundancy, such as RAID 5 and redundant paths to the storage space, whereas the bronze service configuration policy may not require redundant paths or a high RAID level.

[0043] The called switch element configuration 216a, b, c uses (at block 356) interfaces in the lookup service proxy object 254 to query the service attributes of the switch configuration API proxy objects to determine one or more switch configuration API proxy objects capable of configuring switch(s) 132 including paths between the

5 determined storage devices and specified host in a manner that satisfies the called switch element configuration policy parameters. For instance, the gold service configuration policy 202 may require redundant paths through the same or different switches to improve availability, whereas the bronze service configuration policy 200 may not require redundant paths to the storage device.

10 [0044] The called HBA element configuration 218a, b, c uses (at block 358) interfaces in lookup service proxy object 254 to query service attributes for HBA configuration API proxy objects to determine one or more HBA configuration API proxy objects capable of configuring host bus adaptors 234 that can connect to the determined switches and paths that are allocated to satisfy the administrator request.

15 [0045] Note that the above determination of storage devices, switches and host bus adaptors may involve the called device element configuration performing multiple iterations to find some combination of components that can provide the requested storage space allocation to the specified logical volume and host and additionally satisfy the element configuration policy parameters.

20 [0046] After determining the resources 230, 232, and 234 to use to fulfill the administrator UI 's 252 storage allocation request, the called device element configurations 214a, b, c, 216a, b, c, 218a, b, c, and 220a, b, c call the determined configuration APIs to perform the user requested allocation. At block 360, the previously called storage device element configuration 214a, b, c uses the one or more

25 determined storage configuration API proxy objects 224, the APIs therein, to configure the associated storage device(s) to allocate storage space for the requested allocation. At block 364, the switch element configuration 216a, b, c uses the one or more determined switch configuration API proxy objects, and APIs therein, to configure the associated switches to allocate paths for the requested allocation.

[0047] At block 366, the previously called HBA element configuration 218a, b, c uses the determined HBA configuration API proxy objects, and APIs therein, to assign the associated host bus adaptors 234 to the determined path.

[0048] At block 368, the volume manager element configuration 220a, b, c uses the 5 determined volume manager API proxy objects, and APIs therein, to assign the allocated storage space to the logical volumes in the host specified in the administrator UI request.

[0049] The configuration APIs 222, 224, 226, 228, may grant element 10 configurations 214a, b, c, 216a, b, c, 218a, b, c, 220a, b, c access to the API resources on an exclusive or non-exclusive basis according to the lease policy for the configuration API proxy objects.

[0050] The described implementations thus provide a technique to allow for 15 automatic configuration of numerous SAN resources to allocate storage space for a logical volume on a specified host. In the prior art, users would have to select components to assign to an allocation and then separately invoke different configuration tools for each affected component to implement the requested allocation. With the described implementation, the administrator UI or other entity need only specify the new storage allocation one time, and the configuration of the 20 multiple SAN components is performed by singularly invoking one service configuration policy 200, 202, that then invokes the device element configurations.

Using a Defined Service Configuration Policy

to Implement a Resource Allocation

[0051] FIG. 6 illustrates further details of the administrator UI system 252 25 including the lookup service proxy object 254 shown in FIG. 3. The administrator UI 252 further includes a configuration policy tool 270 which comprises a software program that a system administrator may invoke to define and add service configuration policies and allocate storage space to a host bus adaptor (HBA) according to a predefined service configuration policy. A display monitor 272 is 30 attached to the administrator UI 252 to display a graphical user interface (GUI) generated by the configuration policy tool 270.

[0052] FIGs. 7-8 illustrate GUI panels the configuration policy tool 270 displays to allow the administrator UI to operate one of the previously defined service configuration policies to configure and allocated storage space. FIG. 7 is a GUI panel 400 displaying a drop down menu 402 in which the administrator may select one host including one or more bus adaptors (HBA) in the system for which the resource allocation will be made. A descriptive name of the host or any other name, such as the world wide name, may be displayed in the panel drop down menu 402. After selecting a host, the administrator may select from drop down menu 404 a predefined configuration service policy to use to configure the selected host, e.g., bronze, silver, gold, platinum, etc.. Each configuration service policy 200, 202 displayed in the menu 404 has a proxy object 238 registered with the lookup service 250 (FIG. 3). The administrator may obtain more information about the configuration policy parameters for the selected configuration policy displayed in the drop down menu 404 by selecting the “More Info” button 406. The information displayed upon selection of the “More Info” button 406 may be obtained from the service attributes included with the proxy objects 238 for the service configuration policies.

[0053] If the administrator selects one host in drop down menu 402, then the configuration policy tool 270 may determine, according to the logic described below with respect to FIG. 9, those service configuration policies 238 that can be used to configure the selected host, and only display those determined service configuration policies in the drop down menu 404 for selection. Alternatively, the administrator may first select a service configuration policy 200,202 in drop down menu 404, and then the drop down menu 402 would display those hosts that are available to be configured by the selected service configuration policy 200, 202, i.e., those hosts that include a host bus adaptor (HBA) connected to resources, e.g., a switch and storage device, that can satisfy the configuration policy parameters 124 of the elements 106 (FIG. 2), 214a, b, c, 216a, b, c, 218a, b, c, 220a, b, c (FIG. 3), included in the selected service configuration policy.

[0054] After a service configuration policy and host are selected in drop down menus 402 and 404, the administrator may then select the Next button 408 to proceed to the GUI panel 450 displayed in FIG. 8. The panel 450 displays a slider 452 that the

administrator may control to indicate the amount of storage space to allocate to the previously selected host according to the selected service configuration policy. The maximum selectable storage space on the slider 452 is the maximum available for the storage resources that may be configured for the selected host and configuration 5 policy. The minimum storage space indicated on the slider 452 may be the minimum needed to comply with the selected service configuration policy parameters. Panel 450 further displays a text box 454 showing the storage capacity selected on the slider 452. Upon selection of the amount of storage space to allocate using the slider 452 and the Finish button 456, the configuration policy tool 270 would then invoke the 10 selected service configuration policy to allocate the administrator specified storage space using the resources the administrator selected.

[0055] FIGs. 9 and 10 illustrate logic implemented in the configuration policy tool 270 and other of the components in the architecture described with respect to FIGs. 2 and 3 to allocate storage space according to a selected predefined service 15 configuration policy. With respect to FIG. 9, control begins at block 500, where the configuration policy tool 270 is invoked to allocate storage space. The configuration policy tool 270 then determines (at block 502) all the available hosts in the system using the topology database 140 (FIG. 2), 256 (FIG. 3). Alternatively, the configuration policy tool 270 can use the lookup service proxy object 254 to query the 20 service attributes of the proxy objects for the HBA configuration APIs to determine the name of all hosts in the system. A host may include multiple host bus adaptors 234. The name of all the determined hosts are then provided (at block 504) to the drop down menu 402 for administrator selection. The configuration policy tool 270 then displays (at block 506) the panel 400 (FIG. 7) to receive administrator selection 25 of one host and one predefined service configuration policy 200, 202 to use to configure the host.

[0056] Upon receiving (at block 508) administrator selection of one host, the configuration policy tool 270 then queries (at block 510) the service attributes 130 (FIG. 2) of each service configuration policy proxy object 120 (FIG. 2), 238 (FIG. 3) 30 to determine whether the administrator selected host is a member of the service configuration policy, i.e., whether the selected host includes a host bus adaptor (HBA)

arrangement that can satisfy the requirements of the selected service configuration policy 200, 202. As discussed the service attributes 130 of the configuration policy proxy objects 120 (FIG. 2) provide information on all the resources in the system that may be used and configured by the configuration policy. Alternatively, information 5 on the topology of available resources for the host may be obtained by querying the topology database 256, and then a determination can be made as to whether the resources available to the host as indicated in the topology database 256 are capable of satisfying the configuration policy parameters. Still further, a determination can be made of those resources available to the host as indicated in the topology database 256 10 that are also listed in the service attributes 130 of the service configuration policy proxy object 120 indicating resources capable of being configured by the service configuration policy 108 represented by the proxy object. The configuration policy tool 270 then displays (at block 512) the drop down menu 404 with the determined service configuration policies that may be used to configure one host bus adaptor 15 (HBA) 234 in the host selected in drop down menu 402 (FIG. 7)

[0057] Upon receiving (at block 514) administrator selection of the Next button 408 (FIG. 7) with one host and service configuration policy 200, 202 selected, the configuration policy tool 270 then uses the lookup service proxy object 254 to query (at block 518) the service attributes 130 of the selected service configuration policy 20 proxy object 120 (FIG. 2), 238 (FIG. 3) to determine all the host bus adaptors (HBA) available to the selected service configuration policy that are in the selected host and the available storage devices 230 attached to the available host bus adaptors (HBAs) in the selected host. As discussed, such information on the connectedness or topology of the resources is included in the topology database 140 (FIG. 2), 256 (FIG. 25 3). The configuration policy tool 270 then queries (at block 522) the service attributes in the storage device configuration API proxy object 242 to determine the allocatable or available storage space in each of the available storage devices connected to the host subject to the configuration. The total available storage space across all the storage devices is determined (at block 524). The storage space 30 allocated to the host according to the configuration policy may comprise a virtual storage space extending across multiple storage devices. The allocate storage panel

450 (FIG. 8) is then displayed (at block 526) with the slider 452 having as a maximum amount the total storage space in all the available storage devices connected to the host and a minimum amount indicated in the the configuration policy 108, 202 or the configuration policy parameters for the storage device element configuration 214a, b, 5 c (FIG. 3) for the selected configuration policy. Control then proceeds to block 550 in FIG. 10.

[0058] Upon receiving (at block 550) administrator selection of the Next button 456 after administrator selection of an amount of storage space using the slider, the configuration policy tool 270 then determines (at block 552) one available storage 10 device that can provide the administrator selected amount of storage. At block 522, the amount of storage space in each available storage device was determined. The configuration policy tool 270 then queries (at block 554) the service attributes of the selected configuration policy proxy object 238 to determine the available host bus adaptor (HBA) in the selected host that is connected to the determined storage device 15 230 capable of satisfying the administrator selected space allocation. The service attributes are further queried (at block 556) to determine one or more switches in the path between the determined available host bus adaptor (HBA) and the determined storage device. If the selected service configuration policy requires redundant hardware components, then available redundant resources would also be determined. 20 After determining all the resources to use for the allocation that connect to the selected host, the service configuration policy 200, 202 is called (at block 558) to configure the determined resources, e.g., HBA, switch, storage device, and any other components.

[0059] In the above described implementation, the administrator only made one 25 resource selection of a host. Alternatively, the administrator may make additional selections of resources, such as select the host bus adaptor (HBA), switch and/or storage device to use. In such case, upon administrator selection of one additional component to use, the configuration policy tool 270 would determine from the service attributes of the selected service configuration policy the available downstream 30 components that is connected to the previously selected resource instances. Thus,

administrator or automatic selection of an additional component is available for use with a previous administrator selection.

[0060] The above described graphical user interfaces (GUI) allows the administrator to make the minimum necessary selections, such as a host, service configuration 5 policy to use, and storage space to allocate to such host. Based on these selections, the configuration policy tool 270 is able to automatically determine from the registered proxy objects in the look service the resources, e.g., host bus adaptor (HBA), switch, storage, etc., to use to allocate the selected space according to the selected configuration policy without requiring any further information from the 10 administrator. At each step of the selection process, the underlying program components query the system for available resources or options that satisfy the previous administrator selections.

Graphical User Interface to Create

a Service Quality Configuration Policy

[0061] FIGs. 11 and 12 illustrate GUI panels that the configuration policy tool 270 generates to guide the user through creating a service configuration policy, e.g., a bronze quality, gold quality, etc. FIG. 11 illustrates a GUI panel 600 showing all the 20 configurable resources in the system that the user may select to include in a service configuration policy, such as a storage device 230, switch 232, host bus adaptors 234, volume manager program 236, snapshot copy application, backup/archive application, etc. The GUI panel 600 displays graphical icons 602, 604, 606, 608, 610, 612 and associated check boxes 614, 616, 618, 620, 622, and 624. The user may select the 25 resources 602, 604, 606, 608, 610, 612 to add to the defined configuration by selecting the corresponding check boxes 614, 616, 618, 620, 622, 624 associated with the resources. After checking the check boxes corresponding to the resources to add to the service configuration policy, the user may select the Next button 626 to proceed to the GUI panel 650 shown in FIG. 12 to the next step of the service configuration policy definition procedure.

30 [0062] FIG. 12 illustrates a GUI panel 650 in which the user may select an available element configuration 214a, b, c, 216a, b, c, 218a, b, c, 220a, b, c (FIG. 3) for each of

the resources 602, 604, 606, 608, 610, 612 selected in the previous GUI panel 600 in FIG. 11. GUI panel 650 displays icons representing each resource selected from the previous panel 600 and an associated drop down menu 652, 654, 656, 658 for the selected resources. Each drop down menu 652, 654, 656, 658 presents available 5 element configurations 214a, b, c, 216a, b, c, 218a, b, c, 220a, b, c the user may select to add to the service configuration policy. As discussed, the element configurations enable configuration and control over a resource according to a predefined level. For instance, FIG. 12 illustrates the storage device drop down menu 658 showing four different possible predefined storage device element configurations, including the 10 selected configuration. The other drop down menus 654 and 656 show user or default selected element configurations to include in the service configuration policy being defined. The user would use the drop down menus 652, 654, 656, and 658 to select one predefined element configuration for each selected resource to add to the service configuration policy. In this way, the user selects the level of service for 15 each of the selected resources. After the user selects the element configuration to use for each resource, the user may then enter a name for the defined configuration policy in the text box 660 and select a Finish button 662 to create the configuration policy and generate and register a service quality configuration policy proxy object with the lookup service 250.

20 [0063] FIG. 13 illustrates logic implemented in the configuration policy tool 270 and other components in the architecture of FIG. 3 to interact with a user through the GUI panels 600 and 650 to enable a user to create a service configuration policy. Control begins at block 700 with the configuration policy tool 270 being invoked. In response, the configuration policy tool 270 displays the resource select 25 panel 600 (FIG. 11) to allow the user to select displayed resources 602, 604, 606, 608, 610, 612 to add to the service configuration policy being created. In response to receiving (at block 704) user selection of a set of resources and the Next button 626, the configuration policy tool 270 uses the interfaces in the lookup service proxy object 254 to query (at block 706) the service attributes of the element configuration proxy 30 objects 240 (FIG. 3) to determine the name of the element configurations for the user selected resources. The configuration policy tool 270 then displays (at block 708) the

select element configuration panel 650 (FIG. 7) including drop down menus 652, 654, 656, 658 for each user selected resource, where each displayed drop down menu 652, 654, 656, 658 is capable of displaying the names of the element configurations 214a, b, c, 216a, b, c, 218a, b, c, 220a, b, c for the resource, as shown with the storage device drop down menu 658. At block 710, the configuration policy tool 270 detects the selection of the Finish button 660 after the user has selected element configurations in each drop down menu 652, 654, 656, 658 and entered a name for the service configuration policy in the text box 662. At this point, the service configuration policy is defined, and the configuration policy tool 270 generates and registers (at block 712) the user named service configuration policy proxy object 238 with the lookup service 250 (FIG. 3). The configuration policy tool 270 may also generate information on the configuration parameters that would be implemented by the user defined service configuration policy to include in the service attributes for the defined service configuration policy proxy object in the lookup service 250.

15

Additional Implementation Details

[0064] The described implementations presented GUI panels including an arrangement of information and selectable items. Those skilled in the art will appreciate that there are many ways the information and selectable items in the illustrated GUI panels may be aggregated into fewer panels or dispersed across a greater number of panels than shown. Further, additional implementations may provide different layout and user interface mechanisms to allow users to enter the information entered through the discussed GUI panels. In alternative embodiments, users may enter information through line commands as opposed to a GUI.

[0065] The described implementations may be realized as a method, apparatus or article of manufacture using standard programming and/or engineering techniques to produce software, firmware, hardware, or any combination thereof. The term "article of manufacture" as used herein refers to code or logic implemented in hardware logic (e.g., an integrated circuit chip, Field Programmable Gate Array (FPGA), Application Specific Integrated Circuit (ASIC), etc.) or a computer readable medium (e.g., magnetic storage medium (e.g., hard disk drives, floppy disks,, tape, etc.), optical

storage (CD-ROMs, optical disks, etc.), volatile and non-volatile memory devices (e.g., EEPROMs, ROMs, PROMs, RAMs, DRAMs, SRAMs, firmware, programmable logic, etc.). Code in the computer readable medium is accessed and executed by a processor. The code in which preferred embodiments of the

5 configuration discovery tool are implemented may further be accessible through a transmission media or from a file server over a network. In such cases, the article of manufacture in which the code is implemented may comprise a transmission media, such as a network transmission line, wireless transmission media, signals propagating through space, radio waves, infrared signals, etc. Of course, those skilled in the art

10 will recognize that many modifications may be made to this configuration without departing from the scope of the present invention, and that the article of manufacture may comprise any information bearing medium known in the art.

[0066] The implementations were described with respect to the Sun Microsystems, Inc. Jiro network environment that provides distributed computing. However, the

15 described technique for configuration of components may be implemented in alternative network environments where a client downloads an object or code from a server to use to access a service and resources at that server. Moreover, the described configuration policy services and configuration elements that were described as implemented in the Java programming language as Jiro proxy objects may be

20 implemented in any computer architecture known in the art and coded using any known programming language to perform the functions described herein.

[0067] In the described implementations, the storage comprised network storage accessed over a network. Additionally, the configured storage may comprise a storage

25 device directly attached to the host.

[0068] The described logic of FIGs. 4 and 5 concerned a request to add additional storage space to a logical volume. However, the above described architecture and configuration technique may apply to other types of operations involving the allocation of storage resources, such as freeing-up space from one logical volume or

30 requesting a reallocation of storage space from one logical volume to another.

[0069] The configuration policy services 202, 204 may control the configuration elements 214a, b, c, 216a, b, c, 218a, b, c, and 220a, b, c over the Fibre Channel links or use an out-of-band communication channel, such as through a separate LAN connecting the devices 230, 232, and 234.

5 [0070] The configuration elements 214a, b, c, 216a, b, c, 218a, b, c, and 220a, b, c may be located on the same computing device including the requested resource, e.g., storage device 230, switch 232, host bus adaptors 234, or be located at a remote location from the resource being managed and configured.

[0071] In the described implementations, the service configuration policy service

10 configures a switch when allocating storage space to a specified logical volume in a host. Additionally, if there are no switches (fabric) in the path between the specified host and storage device including the allocated space, there would be no configuration operation performed with respect to the switch.

[0072] In the described implementations, the service configuration policy was used

15 to control elements related to the components within a SAN environment. Additionally, the configuration architecture of FIG. 2 may apply to any system in which an operation is performed, such as an allocation of resources, that requires the management and configuration of different resources throughout the system. In such cases, the elements may be associated with any element within the system that is

20 manipulated through a configuration policy service.

[0073] In the described implementations, the architecture was used to alter the allocation of resources in the system. Additionally, the described implementations may be used to control system components through the elements to perform operations other than configuration operations, such as operations managing and

25 controlling the device.

[0074] The above implementations were described with respect to a Fibre Channel environment. Additionally, the above described implementations of the invention may apply to other network environments, such as InfiniBand, Gigabit Ethernet, TCP/IP, the Internet, etc.

30 [0075] In the above described implementations, specific operations were described as being performed by a service configuration policy, device element configuration

and device APIs. Alternatively, functions described as being performed with respect to one type of object may be implemented in another object. For instance, operations described as performed with respect to the element configurations may be performed by the service configuration policies.

5 [0076] The foregoing description of the implementations of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be limited not by this detailed description, but rather by the claims

10 10 appended hereto. The above specification, examples and data provide a complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the spirit and scope of the invention, the invention resides in the claims hereinafter appended.

15

**JIRO, JAVA, SUN, and SUN MICROSYSTEMS are trademarks of Sun Microsystems, Inc. InfiniBand is a service mark of the InfiniBand Trade Association.

WHAT IS CLAIMED IS:

1 1. A method for managing multiple resources in a system, comprising:
2 receiving a user request to generate a configuration policy;
3 receiving user selection of a set of the multiple resources;
4 determining at least one element for each selected resource in the set, wherein
5 each element is capable of managing one of the resources in the system;
6 receiving user selection of one element for each selected resource in the set;
7 and
8 defining the configuration policy to include the user selected elements,
9 wherein invoking the configuration policy further invokes each element defined in the
10 configuration policy to configure the resources associated with the invoked elements.

1 2. The method of claim 1, further comprising:
2 displaying a first user interface enabling the user to select the set of the
3 multiple resources to include in the configuration policy; and
4 displaying a second user interface enabling the user to select the one element
5 for each resource in the set.

1 3. The method of claim 1, wherein the multiple resources include a
2 storage device, a switch, a host adaptor, and a volume manager, wherein the elements
3 managing the storage device allocate the storage space to the host, wherein the
4 elements managing the switch are capable of allocating at least one path in the switch
5 to the storage device to allow the host to access the allocated storage space, wherein
6 the elements managing the host adaptors allocate at least one host adaptor in the host
7 to communicate with the switch to access the allocated storage space, and wherein the
8 elements managing the volume manager assign the allocated storage space in the
9 device to the requested logical volume used by the host.

1 4. The method of claim 3, wherein the system is capable of including
2 multiple storage devices, switches, and host adaptors in the host, and wherein there is
3 at least one separate element to manage each storage device and switch in the system.

1 5. The method of claim 1, wherein each of multiple elements provided for
2 one resource define a different configuration of the resource.

1 6. The method of claim 1, wherein determining the at least one element
2 for each resource comprises:

3 using interfaces in a lookup service proxy object to query element proxy
4 objects to determine a name for each of the element proxy objects.

1 7. The method of claim 6, further comprising:

2 displaying at least one selectable list of the names of each of the element proxy
3 objects for each resource, wherein the user selects one element for each resource from
4 the selectable lists.

1 8. A method for configuring multiple resources in the system,
2 comprising:

3 receiving user selection of one of multiple configuration policies, wherein
4 each configuration policy defines resources to configure and one element for each
5 resource to configure, wherein each element specifies configuration parameters to use
6 to configure the resource;

7 receiving user selection of an instance of one resource to configure, wherein
8 the user selected resource instance is capable of being configured by the configuration
9 policy;

10 determining additional resource instances that are configured by the selected
11 configuration policy; and

12 calling the elements defined for the selected configuration policy to configure
13 the user selected resource instance and the determined additional resource instances
14 according to the element configuration parameters.

1 9. The method of claim 8, further comprising:
2 displaying a first interface listing the multiple configuration policies, wherein
3 the user selects one configuration policy from the list; and
4 displaying at a second interface enabling the user to select the instance of the
5 resource to configure.

1 10. The method of claim 8, further comprising:
2 querying information indicating the resource instances available for the
3 configuration, wherein the information indicates the connectedness of the resource
4 instances, wherein the determined additional resource instances are connected.

1 11. The method of claim 8, wherein one of the resources to configure
2 comprises a storage device further comprising:
3 determining available storage space at each storage device instance that is
4 available to the user selected resource instance;
5 receiving user selection of an amount of storage space to allocate; and
6 determining one storage device instance including the user selected amount of
7 storage space, wherein calling the elements to configure each user selected resource
8 further comprises calling a storage element to configure the determined storage device
9 instance to allocate the user selected amount of storage space to the configuration.

1 12. The method of claim 11, further comprising:
2 displaying a storage allocation interface displaying the available storage space,
3 wherein the user enters the selected amount of storage space through the displayed
4 allocation interface.

1 13. The method of claim 8, wherein the multiple resources include a
2 storage device and a host adaptor, and wherein the user selected resource comprises a
3 host including at least one host adaptor, and wherein the determined additional
4 resources instances comprise one instance of the host adaptor and storage device
5 resources.

1 14. The method of claim 13, wherein the multiple resources further include
2 a switch, and wherein the determined additional resources instances further include
3 one instance of the switch resource.

1 15. The method of claim 13, wherein determining the instance of the host
2 adaptor and storage device comprises querying information indicating host adaptor
3 and storage device instances capable of being configured according to the
4 configuration parameters and the topology of the host adaptor and storage device
5 instances, and wherein the determined host adaptor and storage device instances to
6 use in the configuration are connected.

1 16. The method of claim 15, wherein the multiple resources further include
2 a switch, and wherein the determined additional resources instances further include
3 one instance of the switch resource, wherein the determined switch instance is in a
4 path between the determined host adaptor and storage device instances.

1 17. The method of claim 16, wherein the element managing the storage
2 device allocates storage space to the host, wherein the element managing the switch is
3 capable of allocating at least one path in the switch to the storage device to allow the
4 host to access the allocated storage space, wherein the element managing the host
5 adaptors allocates at least one host adaptor in the host to communicate with the switch
6 to access the allocated storage space.

1 18. The method of claim 8, wherein each of multiple elements provided for
2 one resource define a different configuration of the resource.

1 19. The method of claim 8, further comprising:
2 querying configuration policy proxy objects in a lookup service to determine
3 configuration policies;

4 displaying a user interface listing the determined configuration policies,
5 wherein the user selects one of the configuration policies from the list;
6 downloading the configuration policy proxy object for the selected
7 configuration policy from the lookup service; and
8 using an interface in the downloaded configuration policy proxy object to call
9 the elements for each resource to configure the user selected and additional resource
10 instances according to the element configuration.

1 20. The method of claim 8, wherein determining the additional instances of
2 the resource further comprises:
3 querying attributes associated with a proxy object in a lookup service for the
4 user selected configuration policy to determine resource instances capable of being
5 configured by the selected configuration policy.

1 21. A system for managing multiple resources, comprising:
2 a computer readable medium including at least one element for each of the
3 managed resources in the system, wherein each element is capable of managing one of
4 the resources in the system;
5 means for receiving a user request to generate a configuration policy;
6 means for receiving user selection of a set of the multiple resources;
7 means for determining at least one element for each selected resource in the
8 set;
9 means for receiving user selection of one element for each selected resource in
10 the set; and
11 means for defining the configuration policy to include the user selected
12 elements, wherein invoking the configuration policy further invokes each element
13 defined in the configuration policy to configure the resources associated with the
14 invoked elements.

1 22. The system of claim 21, further comprising:
2 means for displaying a first user interface enabling the user to select the set of
3 the multiple resources to include in the configuration policy; and
4 means for displaying a second user interface enabling the user to select the one
5 element for each resource in the set.

1 23. The system of claim 21, wherein the multiple resources include a
2 storage device, a switch, a host adaptor, and a volume manager, wherein the elements
3 managing the storage device allocate the storage space to the host, wherein the
4 elements managing the switch are capable of allocating at least one path in the switch
5 to the storage device to allow the host to access the allocated storage space, wherein
6 the elements managing the host adaptors allocate at least one host adaptor in the host
7 to communicate with the switch to access the allocated storage space, and wherein the
8 elements managing the volume manager assign the allocated storage space in the
9 device to the requested logical volume used by the host.

1 24. The system of claim 31, wherein the managed resources are capable of
2 including multiple storage devices, switches, and host adaptors in the host, and,
3 wherein there is at least one separate element in the computer readable medium to
4 manage each storage device and switch in the system.

1 25. The system of claim 21, wherein each of multiple elements provided
2 for one resource define a different configuration of the resource.

1 26. The system of claim 21, wherein the computer readable medium
2 further includes element proxy objects and a lookup service proxy object, and wherein
3 the means for determining the at least one element for each resource performs:
4 using interfaces in the lookup service proxy object to query element proxy
5 objects to determine a name for each of the element proxy objects.

1 27. The system of claim 26, further comprising:
2 means for displaying at least one selectable list of the names of each of the
3 element proxy objects for each resource, wherein the user selects one element for each
4 resource from the selectable lists.

1 28. A system for configuring multiple resources, comprising:
2 a computer readable medium including:
3 (i) at least one element for each of the managed resources in the
4 system, wherein each element is capable of managing one of the resources in
5 the system, and wherein each element specifies configuration parameters to
6 use to configure the resource;
7 (ii) configuration policies, wherein each configuration policy defines
8 resources to configure and one element for each resource to configure;
9 means for receiving user selection of one of the configuration policies,
10 means for receiving user selection of an instance of one resource to configure,
11 wherein the user selected resource instance is capable of being configured by the
12 configuration policy;
13 means for determining additional resource instances that are configured by the
14 selected configuration policy; and
15 means for calling the elements defined for the selected configuration policy to
16 configure the user selected resource instance and the determined additional resource
17 instances according to the element configuration parameters.

1 29. The system of claim 28, further comprising:
2 means for displaying a first interface listing the multiple configuration
3 policies, wherein the user selects one configuration policy from the list; and
4 means for displaying at a second interface enabling the user to select the
5 instance of the resource to configure.

1 30. The system of claim 28, further comprising:
2 means for querying information indicating the resource instances available for
3 the configuration, wherein the information indicates the connectedness of the resource
4 instances, wherein the determined additional resource instances are connected.

1 31. The system of claim 28, wherein one of the resources to configure
2 comprises a storage device further comprising:
3 means for determining available storage space at each storage device instance
4 that is available to the user selected resource instance;
5 means for receiving user selection of an amount of storage space to allocate;
6 and
7 means for determining one storage device instance including the user selected
8 amount of storage space, wherein the means for calling the elements to configure each
9 user selected resource further performs calling a storage element to configure the
10 determined storage device instance to allocate the user selected amount of storage
11 space to the configuration.

1 32. The system of claim 31, further comprising:
2 means for displaying a storage allocation interface displaying the available
3 storage space, wherein the user enters the selected amount of storage space through
4 the displayed allocation interface.

1 33. The system of claim 28, wherein the multiple resources include a
2 storage device and a host adaptor, and wherein the user selected resource comprises a
3 host including at least one host adaptor, and wherein the determined additional
4 resources instances comprise one instance of the host adaptor and storage device
5 resources.

1 34. The system of claim 33, wherein the multiple resources further include
2 a switch, and wherein the determined additional resources instances further include
3 one instance of the switch resource.

1 35. The system of claim 33, wherein the means for determining the
2 instance of the host adaptor and storage device queries information indicating host
3 adaptor and storage device instances capable of being configured according to the
4 configuration parameters and the topology of the host adaptor and storage device
5 instances, and wherein the determined host adaptor and storage device instances to
6 use in the configuration are connected.

1 36. The system of claim 35, wherein the multiple resources further include
2 a switch, and wherein the determined additional resources instances further include
3 one instance of the switch resource, wherein the determined switch instance is in a
4 path between the determined host adaptor and storage device instances.

1 37. The system of claim 36, wherein the element managing the storage
2 device allocates storage space to the host, wherein the element managing the switch is
3 capable of allocating at least one path in the switch to the storage device to allow the
4 host to access the allocated storage space, wherein the element managing the host
5 adaptors allocates at least one host adaptor in the host to communicate with the switch
6 to access the allocated storage space.

1 38. The system of claim 28, wherein each of multiple elements provided
2 for one resource define a different configuration of the resource.

1 39. The system of claim 28, further comprising:
2 means for querying configuration policy proxy objects in a lookup service to
3 determine configuration policies;
4 means for displaying a user interface listing the determined configuration
5 policies, wherein the user selects one of the configuration policies from the list;
6 means for downloading the configuration policy proxy object for the selected
7 configuration policy from the lookup service; and
8 means for using an interface in the downloaded configuration policy proxy
9 object to call the elements for each resource to configure the user selected and
10 additional resource instances according to the element configuration.

1 40. The system of claim 28, wherein the means for determining the
2 additional instances of the resource further performs:

3 querying attributes associated with a proxy object in a lookup service for the
4 user selected configuration policy to determine resource instances capable of being
5 configured by the selected configuration policy.

1 41. An article of manufacture including code for managing multiple
2 resources in a system by:

3 receiving a user request to generate a configuration policy;

4 receiving user selection of a set of the multiple resources;

5 determining at least one element for each selected resource in the set, wherein
6 each element is capable of managing one of the resources in the system;

7 receiving user selection of one element for each selected resource in the set;

8 and

9 defining the configuration policy to include the user selected elements,
10 wherein invoking the configuration policy further invokes each element defined in the
11 configuration policy to configure the resources associated with the invoked elements.

1 42. The article of manufacture of claim 41, further comprising:

2 displaying a first user interface enabling the user to select the set of the
3 multiple resources to include in the configuration policy; and

4 displaying a second user interface enabling the user to select the one element
5 for each resource in the set.

1 43. The article of manufacture of claim 41, wherein the multiple resources
2 include a storage device, a switch, a host adaptor, and a volume manager, wherein the
3 elements managing the storage device allocate the storage space to the host, wherein
4 the elements managing the switch are capable of allocating at least one path in the
5 switch to the storage device to allow the host to access the allocated storage space,
6 wherein the elements managing the host adaptors allocate at least one host adaptor in
7 the host to communicate with the switch to access the allocated storage space, and

8 wherein the elements managing the volume manager assign the allocated storage
9 space in the device to the requested logical volume used by the host.

1 44. The article of manufacture of claim 43, wherein the system is capable
2 of including multiple storage devices, switches, and host adaptors in the host, and
3 wherein there is at least one separate element to manage each storage device and
4 switch in the system.

1 45. The article of manufacture of claim 41, wherein each of multiple
2 elements provided for one resource define a different configuration of the resource.

1 46. The article of manufacture of claim 41, wherein determining the at
2 least one element for each resource comprises:
3 using interfaces in a lookup service proxy object to query element proxy
4 objects to determine a name for each of the element proxy objects.

1 47. The article of manufacture of claim 46, further comprising:
2 displaying at least one selectable list of the names of each of the element proxy
3 objects for each resource, wherein the user selects one element for each resource from
4 the selectable lists.

1 48. An article of manufacture method for configuring multiple resources in
2 the system by:
3 receiving user selection of one of multiple configuration policies, wherein
4 each configuration policy defines resources to configure and one element for each
5 resource to configure, wherein each element specifies configuration parameters to use
6 to configure the resource;
7 receiving user selection of an instance of one resource to configure, wherein
8 the user selected resource instance is capable of being configured by the configuration
9 policy;

10 determining additional resource instances that are configured by the selected
11 configuration policy; and
12 calling the elements defined for the selected configuration policy to configure
13 the user selected resource instance and the determined additional resource instances
14 according to the element configuration parameters.

1 49. The article of manufacture of claim 48, further comprising:
2 displaying a first interface listing the multiple configuration policies, wherein
3 the user selects one configuration policy from the list; and
4 displaying at a second interface enabling the user to select the instance of the
5 resource to configure.

1 50. The article of manufacture of claim 48, further comprising:
2 querying information indicating the resource instances available for the
3 configuration, wherein the information indicates the connectedness of the resource
4 instances, wherein the determined additional resource instances are connected.

1 51. The article of manufacture of claim 48, wherein one of the resources to
2 configure comprises a storage device further comprising:
3 determining available storage space at each storage device instance that is
4 available to the user selected resource instance;
5 receiving user selection of an amount of storage space to allocate; and
6 determining one storage device instance including the user selected amount of
7 storage space, wherein calling the elements to configure each user selected resource
8 further comprises calling a storage element to configure the determined storage device
9 instance to allocate the user selected amount of storage space to the configuration.

1 52. The article of manufacture of claim 51, further comprising:
2 displaying a storage allocation interface displaying the available storage space,
3 wherein the user enters the selected amount of storage space through the displayed
4 allocation interface.

1 53. The article of manufacture of claim 48, wherein the multiple resources
2 include a storage device and a host adaptor, and wherein the user selected resource
3 comprises a host including at least one host adaptor, and wherein the determined
4 additional resources instances comprise one instance of the host adaptor and storage
5 device resources.

1 54. The article of manufacture of claim 53, wherein the multiple resources
2 further include a switch, and wherein the determined additional resources instances
3 further include one instance of the switch resource.

1 55. The article of manufacture of claim 53, wherein determining the
2 instance of the host adaptor and storage device comprises querying information
3 indicating host adaptor and storage device instances capable of being configured
4 according to the configuration parameters and the topology of the host adaptor and
5 storage device instances, and wherein the determined host adaptor and storage device
6 instances to use in the configuration are connected.

1 56. The article of manufacture of claim 55, wherein the multiple resources
2 further include a switch, and wherein the determined additional resources instances
3 further include one instance of the switch resource, wherein the determined switch
4 instance is in a path between the determined host adaptor and storage device
5 instances.

1 57. The article of manufacture of claim 56, wherein the element managing
2 the storage device allocates storage space to the host, wherein the element managing
3 the switch is capable of allocating at least one path in the switch to the storage device
4 to allow the host to access the allocated storage space, wherein the element managing
5 the host adaptors allocates at least one host adaptor in the host to communicate with
6 the switch to access the allocated storage space.

1 58. The article of manufacture of claim 48, wherein each of multiple
2 elements provided for one resource define a different configuration of the resource.

1 59. The article of manufacture of claim 48, further comprising:
2 querying configuration policy proxy objects in a lookup service to determine
3 configuration policies;
4 displaying a user interface listing the determined configuration policies,
5 wherein the user selects one of the configuration policies from the list;
6 downloading the configuration policy proxy object for the selected
7 configuration policy from the lookup service; and
8 using an interface in the downloaded configuration policy proxy object to call
9 the elements for each resource to configure the user selected and additional resource
10 instances according to the element configuration.

1 60. The article of manufacture of claim 48, wherein determining the
2 additional instances of the resource further comprises:
3 querying attributes associated with a proxy object in a lookup service for the
4 user selected configuration policy to determine resource instances capable of being
5 configured by the selected configuration policy.

1/13





FIG. 1



EIGI 2

3/13




FIG. 3

4/13




FIG. 4

5/13




FIG. 5

6/13




FIG. 6

7/13




FIG. 7

8/13




FIG. 8

9/13




FIG. 9

10/13




FIG. 10

11/13




FIG. 11

12/13




FIG. 12

13/13




FIG. 13